

RAND



Contractual Component Repair Policy

*A Key to Improving
Depot Responsiveness*

Mary E. Chenoweth, John B. Abell

19950302 075

Project AIR FORCE



The research reported here was sponsored by the United States Air Force under Contract F49620-91-C-0003. Further information may be obtained from the Strategic Planning Division, Directorate of Plans, Hq USAF.

Library of Congress Cataloging in Publication Data

Chenoweth, Mary E.

Contractual component repair policy : a key to improving depot responsiveness / Mary E. Chenoweth, John B. Abell.

p. cm.

"Prepared for the United States Air Force."

Includes bibliographical references.

"MR-440-AF."

ISBN 0-8330-1607-5

1. Airplanes, Military—United States—Maintenance and repair—Management. 2. Defense contracts—United States—Management.

I. Abell, John B. II. United States. Air Force. III. Title.

UG1243.C44 1994

358.4'183—dc20

94-42951

CIP

RAND
Copyright © 1994

RAND is a nonprofit institution that helps improve public policy through research and analysis. RAND's publications do not necessarily reflect the opinions or policies of its research sponsors.

Published 1994 by RAND

1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138

To order RAND documents or to obtain additional information, contact Distribution Services: Telephone: (310) 451-7002; Fax: (310) 451-6915; Internet: order@rand.org.

RAND

Contractual Component Repair Policy

*A Key to Improving
Depot Responsiveness*

Mary E. Chenoweth, John B. Abell

*Prepared for the
United States Air Force*

Project AIR FORCE

Preface

This report addresses issues in contract maintenance policy. Specifically, it focuses on the Air Force's management of contractual repair of aircraft recoverable components. This research has its roots in the Uncertainty Project, a major project that began in late 1984 and led to a major body of work in enhancing the responsiveness of depot-level component repair (Cohen et al., 1991). The intuition supporting many of the hypotheses developed in this research came from preceding research on issues affecting the responsiveness of depot-level component repair.

This work was carried out in the Logistics Project of Project AIR FORCE, RAND's federally funded research and development center (FFRDC) funded by the United States Air Force. It was sponsored by the Deputy Chief of Staff/Logistics, Headquarters, USAF. It should be of interest to logistics managers and analysts throughout the Air Force logistics system, especially those involved in depot-level management and policymaking, and to logisticians in the other military departments and in the Office of the Secretary of Defense.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution	
Availability Codes	
Dist	Avail and/or Special
A-1	

Contents

Preface	iii
Figures	vii
Summary	ix
Acknowledgments	xv
Glossary	xvii
1. INTRODUCTION	1
Research Purpose	2
Research Approach	3
What Follows	5
2. WHAT GOES ON WITH CONTRACTUAL REPAIR?	6
Contract Formulation	6
Determining Repair Requirements	6
Potential Contractors	8
Contract Negotiations	9
Source Selection	11
Workload Formulation	11
Determining Quarterly Repair Requirements	11
Workload Allocation	12
Executing a Delivery Order for Component Repair	12
Repair Process	13
Unserviceable Arrival at the Contractor	13
Induction into Repair	13
Repair	14
Shipment from the Contractor	15
3. HOW RESPONSIVE IS CONTRACTUAL REPAIR?	16
Repair Flow Times	17
Mean Repair Flow Times Do Not Tell the Full Story	21
Important Implications of Uncertainty in Repair Requirements	22
Perceptions of Contractual Repair Responsiveness	23
4. POLICY DIRECTIONS FOR ENHANCING THE RESPONSIVENESS OF CONTRACTUAL REPAIR	25
Improve Organizational Arrangements for Contractual Repair Management	25
Improve the Responsiveness of Contract Repair	27
Reduce AWP Times at Contractual Repair Facilities	28
Ship Assets Directly Between Bases and Contractors	29
Add Responsiveness Incentives to Repair Contracts	31
Use Responsiveness as a Contractor Selection Criterion	32
Extend the Use of DRIVE to Contractual Repair	32

Improve Requirements Determination and Workload Allocation to Reduce Uncertainty	34
Stabilize Repair Requirements and Funding	35
Allocate the Uncertain Portion of the Workload for Split Items to Organic Sources of Repair	35
5. SUMMARY OF RECOMMENDATIONS	38
Appendix	
A. CONTRACTOR REPAIR PERFORMANCE INFORMATION SYSTEMS.....	41
B. STATISTICAL LIMITATIONS OF RESPONSIVENESS DATA.....	43
C. CONTRACTOR RESPONSIVENESS DATA ANALYSIS	44
Bibliography	47

Figures

S.1. Repair Flow Times for Items with 30-Day Negotiated Flow Times	xii
3.1. Repair Flow Times for 19 Contracts with Initial Reporting Beginning in January 1993	19
3.2. Mean Repair Flow Times for 19 Contracts with Initial Reporting Beginning in January 1993	19
3.3. Repair Flow Times for Items with 30-Day Negotiated Flow Times	21
3.4. Repair Flow Times for Items with 45-Day Negotiated Flow Times	21
4.1. Organizational Chart of Key Contract Maintenance Decisionmakers	26
4.2. Expenditures by Source Type: 1980-1988	36

Summary

Contractual depot-level maintenance involves the repair of recoverable spares, engine repair and overhaul, and programmed maintenance of weapon systems by private companies. As recently as 1992, OSD estimated that the Air Force spent \$1.2 billion of its total \$3.7 billion for depot-level maintenance in the private sector (Kingsbury, 1992). Today the Air Force spends about 40 percent of its depot-level budget on contractor support.¹

This report documents an initial examination of the responsiveness of contractual repair of depot-level recoverables, in particular avionics components managed by the Ogden Air Logistics Center (ALC).² Short repair cycle times help enhance depot repair responsiveness in two ways: by reducing the investment required by the Air Force to maintain a given level of support and by enabling logistics organizations to respond more readily to urgent, unanticipated demands.

Although this research originally intended to explore the feasibility of extending a methodology to set repair and distribution priorities, called Distribution and Repair in Variable Environments (DRIVE) to contractual repair, it quickly became clear that prioritization alone would not achieve reasonable levels of responsiveness. Other, more fundamental problems would need to be addressed. As a result, this report:

- Outlines the major steps involved in managing and executing contractual component repair
- Assesses contractual repair responsiveness in terms of repair flow times for a select group of components
- Suggests directions that the Air Force might take and hypotheses that it might evaluate that promise enhanced contractual repair responsiveness.

¹Communications with the United States Air Force, Office of the Deputy Chief of Staff, Logistics, May 2, 1994.

²It focuses on the contractual repair of recoverable spares, also known as component repair, for aircraft systems only. In the logistician's lexicon, an unserviceable component arrives at a depot because the base believes it to be faulty or because the base could not restore it to serviceable condition locally. This repairable component becomes a serviceable asset once it successfully emerges from depot repair.

We based our findings on a series of interviews with key players at HQ AFMC, Ogden ALC, and two of its larger contractors: Lockheed, Fort Worth Division (formerly General Dynamics) and Westinghouse Electric Corporation. We consulted two recommended data sources: the Government-Furnished Material Transaction Reporting System (G009), indicating repair flow times or actual elapsed times for critical repair process segments; and the Contract Depot Maintenance Production and Costs System (G072D), supplying repair costs.³ Because of data quality problems that were resolved by January 1993, we limited our analysis to the activity under 19 contracts during the last three quarters of FY1993 (January 1993 to October 1993).

What Goes On with Contractual Repair?

One way to describe contractual repair is as a series of three primary processes: establishing the contract, allocating the repair workload, and executing the repair. Key decisions and policies apply to each. Of the many significant factors involved in the evolution of a contract, we describe four: the estimation of repair requirements over the life of the contract, the search for potential contractors, the contractual negotiation, and source selection. Clearly, the government seeks contractors able to produce a quality product for a reasonable price. But the terms and conditions the contractor must operate under has much to do with his responsiveness.

The uncertainties involved in the eventual issuance of firm delivery orders together with material support policies that are largely reactive hamper responsiveness. The allocation of repair workload for a particular quarter often depends on commodity trend projections that are updated by the item manager (IM) and production management specialist (PMS) at the ALC. A delivery order gives the contractor legal authorization to begin work once he receives the unserviceable item.⁴ The typical contract stipulates a mean repair flow time on the basis of the calendar time that the contractor may take to complete a repair. Ogden does not count the time an asset spends awaiting parts (AWP) that are not available at the contractor's facility. During the time a contractor waits for government-furnished material (GFM) or other ordered material to arrive, the "clock" stops. It starts again once he receives the needed parts. Contractors

³These information systems were developed for different purposes than ours, i.e., for assessing contractual repair responsiveness. Even so, they provided detailed information pertaining to three discrete segments: waiting for induction, in repair, and waiting for shipment.

⁴Although not grammatically correct, throughout the Air Force logistics system the words "repairable" and "serviceable" denote the component or asset in one of its conditions while in the repair pipeline.

often repair major components called line-replaceable units (LRUs) by repairing their constituent components, called shop-replaceable units (SRUs), as needed. The contract often requires that the contractor ship the serviceable back to the ALC for subsequent shipment to a base.

ALC databases capture three primary repair segments: waiting for induction into the repair process, in the repair process, and waiting for shipment as a serviceable.

How Responsive Is Contractual Repair?

In the 19 contracts that we analyzed, we found very high actual repair flow times. We divided our sample of repaired items into those with *negotiated flow times* of 30 and 45 days. These times exclude any time when repair is not possible (AWP). Items with negotiated flow times of 30 days achieved, on average, actual repair flow times of 59 days. Items with negotiated flow times of 45 days achieved, on average, actual repair flow times of 77 days. (The actual flow times include some AWP time that we were unable to quantify.) The Coronet Deuce demonstration at Ogden ALC demonstrated that organic facilities repairing very similar items can achieve actual repair flow times roughly an order of magnitude smaller than these.

These flow times look even longer when we consider absolute repair flow times instead of averages. As Figure S.1 shows, for items with negotiated flow times of 30 days, only 30 percent had actual flow times of 30 days or less, but almost 60 percent of the items with the lowest flow times had an average of 30 days. A significant fraction have actual flow times of 90 days or more. The analogous outcome for items with negotiated flow times of 45 days is even more striking.

We found no strong correlation between repair costs and flow times. If expensive items cost more to repair, this lack of correlation suggests that the system may not give higher priority to expensive assets than to cheaper ones. If true, this observation has policy and pipeline-cost implications. We did observe some correlation between GFM costs and flow times for assets with 45-day requirements. This suggests that assets with higher GFM costs tend to have longer repair flow times, not an unreasonable hypothesis given current material support policies.




Figure S.1—Repair Flow Times for Items with 30-Day Negotiated Flow Times

Suggestions for Improving the Responsiveness of Contractual Component Repair

Some office within the Air Force Materiel Command (AFMC) and each ALC should have sufficient oversight authority to initiate policy changes that will lead to greater contractor responsiveness. That office would represent the buyer of contractual services. It needs to have the authority and incentives to initiate improvements to contractual repair policy that will enhance repair responsiveness and increase its cost effectiveness.

Emphasizing the value of contractor responsiveness throughout all aspects of contractual component repair management is an important step toward achieving shorter flow times. Developing an objective, accurate means of measuring contractor performance provides many benefits, not just in awarding contracts, but also in their day-to-day management. Incorporating repair flow times, both with and without AWP times, as one dimension of depot-level costs gives decisionmakers a tool they need. They cannot just cut costs without taking responsiveness into consideration and the only way to do that is to measure it. We suggest that modest changes to G009 reporting requirements would provide such information on an important class of contracts. The emphasis on responsiveness must also extend to contractual incentives. Requirements contracts fail to provide the right ones and result in higher total costs when

tradeoffs between repair cycle times and spares investments are explicitly considered.

Rethinking material support policies addresses a critical obstacle to enhanced responsiveness. Since the contractor relies on the Military Standard Requisitioning and Issue Procedures (MILSTRIP) system and generally does not lay in adequate stocks of repair parts, the Air Force does not count the time that elapses while the contractor waits for requisitioned government-furnished material. Thus, a contractor's performance is not affected by AWP time and he has little motivation to raise material support issues with the Air Force. Given the fact that the government currently manages many of these parts, the approach is fair. But the Air Force has opportunities to reduce total repair flow times through material support policies that lead to fewer AWP incidents. A policy that calls for adequate rotatable spares or on-hand serviceable SRUs could reduce the time an LRU must wait during the repair process, thus improving responsiveness.

Improving the process by which the ALCs generate contract repair requirements used in negotiating repair contracts would help reduce the amount of uncertainty involved with anticipated repair quantities, which is necessary to reduce unit repair costs. Once the ALCs begin measuring repair flow times, they can take the next step in determining the value of better responsiveness in terms of total costs to the depot. Many people at the ALCs said that great uncertainty over projected repair requirements undercuts the government's ability to negotiate more favorable terms. Linking the repair estimation process to aircraft availability goals could help in projecting better repair mixes and quantities.

Eliminating secondary shipments to the fullest extent possible would vastly reduce transportation times and increase depot-level responsiveness to the bases. Retrograde and resupply delays increase with every unnecessary shipment of an item from a base or contractor to an ALC. Each time an item passes through an ALC, time in the pipeline increases. Moreover, the ALC can distribute items more effectively if it updates an item's destination as close to the end of the repair process as possible. AFMC could require that all contracts conform to this policy. Such an approach would lead to immediate benefits in getting assets to the customer faster. Air Force bases need to know where to send their unserviceables and the contractor needs to know where to ship the serviceables. Both decisions could be made in near-real-time, provided some changes in key information systems and transportation policies were made.

Evaluating the effectiveness and total costs of moving to a policy that allocates the uncertainty in predicting repairable generations to the organic source of

repair for split items (those repaired by both an ALC and a contractor) could offer new opportunities for reducing contractual repair uncertainties and repair costs.

Extending the use of DRIVE to contractual component repair for repair requirements determination, recoverable asset management, and serviceable asset allocation may have significant payoff. This methodology could support all of these processes. In situations where there is sufficient scope of repair, DRIVE's application can be extended to repair prioritization. In situations where the contractor repairs both LRUs and their indentured recoverable SRUs, DRIVE's application can be extended to proactive SRU repair.

Acknowledgments

A series of interviews with individuals within the Air Force and the private sector provided the fundamental basis of the findings discussed in this report. We spent hours, in some cases days, with individuals who painstakingly took us through their operations and decisionmaking processes. These individuals include:

William Ferris, Lockheed Aircraft Corporation, Fort Worth Division; John G. Graves, Aerospace Logistics Support Division, Westinghouse Electric Corporation; William Ernst, Director of Contracting (Retired), OO-ALC/PK; Cathy Heywood, OO-ALC/FMIC; Maurice Carter, OO-ALC/FM; David Stevens, OO-ALC/LAAR, and respective production management specialists within his group; Vernon Williams, OO-ALC/PKQ, and many very helpful contracting officers; Kennon Cooksey, AFMC/LGPW; Victor J. Presutti, Jr., AFMC/XPS; Curtis Neumann, AFMC/XPSA; Robert McCormick, AFMC/XPSA; and Ann Fee, AFMC/FMI. We are grateful to others, countless in number, who provided us information necessary to this research. We thank those individuals at HQ AFMC; Ogden ALC; San Antonio ALC, G009 OPR; and Sacramento ALC, G072D OPR.

We are also indebted to our colleagues at RAND, Frank Camm, Lionel Galway, Timothy Bonds, and Laura Zakaras, for their thorough reviews of the manuscript and helpful comments and suggestions.

Glossary

AFLC	Air Force Logistics Command (now AFMC)
AFMC	Air Force Materiel Command
ALC	Air Logistics Center
AWP	Awaiting Parts
CAP	Contractor Acquired Property
CPARS	Contractor Performance Assessment Reporting System
DBOF	Defense Business Operations Fund
DCS	Deputy Chief of Staff
DLR	Depot-Level Recoverable
DMRD	Defense Management Report Decision
DoD	Department of Defense
DRIVE	Distribution and Repair in Variable Environments
FAR	Federal Acquisition Regulation
FFRDC	Federally Funded Research and Development Center
FM	Financial Management
GFM	Government Furnished Material
ICS	Interim Contractor Support
IM	Item Manager
IPD	Integrated Product Development
IWSM	Integrated Weapon System Management
LG	DCS/Logistics
LRU	Line-Replaceable Unit
MA	Maintenance
MICAP	Mission Capability
MILSTRIP	Military Standard Requisitioning and Issue Procedures
MISTR	Management of Items Subject to Repair
MM	DCS/Materiel Management
NFT	Negotiated Flow Time
NSN	National Stock Number
OEM	Original Equipment Manufacturer
OFPP	Office of Federal Procurement Policy
OO	Ogden Air Logistics Center
OSD	Office of the Secretary of Defense
PK	DCS/Contracting
PMS	Production Management Specialist
R&D	Research and Development

Recoverable	An asset subject to repair, in contrast to a consumable
Repairable	An unserviceable asset that can be economically repaired
Rotable Spares	Serviceable, on-hand SRUs
SAF	Secretary of the Air Force
SPD	System Program Director
SPO	System Program Office
SRU	Shop-Replaceable Unit
SSM	System Support Manager
USAF	United States Air Force
VRS	Vendor Rating System
XR	DCS/Requirements
UMMIPS	Uniform Materiel Movement and Issue Priority System

1. Introduction

Contract maintenance accounts for a large and growing share of the business the Air Force conducts at the depot level. In fact, contractual services accounted for almost one-third of Air Force expenditures on depot-level maintenance in 1992—\$1.2 billion of a total \$3.7 billion (Kingsbury, 1992). They reportedly represent an even larger share today, summing to about 40 percent of depot-level maintenance expenditures.¹ Contractual sources of repair represent an integral part of the logistics operation. The Air Force allocates repair workload to contractors especially heavily in the early life of a weapon system. Some component repair always remains with contractors, i.e., the Air Force never develops organic capability for some assets.² Thus, it is fair to say that virtually every Air Force weapon system has components that are repaired contractually.

The Air Force is under intense pressure to improve the performance of its depot-level maintenance. Defense Management Report Decision (DMRD) 908 has already dictated that the Air Force save \$1.1 billion over FYs 1991 to 1995 and has provided the basis for removing that money from the Air Force budget to induce change (Kingsbury, 1992). Efforts to downsize the support structure will continue to tighten the budget. Improvements in contractor performance can help the Air Force respond to this budgetary pressure in two ways. They can contribute directly in the 40 percent of maintenance provided by contractors, and they can ensure a viable alternative to organic sources and motivate continuing improvement in the organic provision of maintenance services.

The Department of Defense (DoD) spurred on its objective of introducing new financial incentives when it combined each service's supply and repair accounts, their stock and industrial funds, respectively, into one revolving fund called the Defense Business Operating Fund. A related change, called the stock funding of depot-level recoverables, requires that each service charge the user the full cost of depot-level support.³ This approach works on the principle that by charging the

¹Communications with the United States Air Force, Office of the Deputy Chief of Staff, Logistics, May 2, 1994.

²This is especially true of weapon systems with small force sizes and of other components with high capital investment costs for establishing organic capability.

³The Air Logistics Centers (ALCs) previously issued recoverable assets at no cost to the requisitioner, but today a financial transaction takes place.

base the full cost of depot-level support, more bases will choose to make repairs at intermediate and organizational levels, i.e., closer to the aircraft. In combination, these changes create a stronger demand from the depot's customers for better service, including the service provided by the private sector.

As to putting the customer first, at least one practice corporate America has learned is that delivering services ever faster makes for good business and happier customers. The Air Force has already embarked on the road to improving the responsiveness of its organic repair operations with its development of DRIVE (Distribution and Repair in Variable Environments), its demonstration of the two-level-maintenance concept for avionics components, and its exploration of a concept called *lean logistics*, as well as other initiatives. All are aimed at making the depot more responsive to the customer or aircraft. Because of the widespread use of contractual sources of repair, a similar spotlight should be cast in the direction of how the Air Force uses its private-sector sources.

We have a definitive notion of responsiveness. It includes relevance, timeliness, and robustness. A repair system is relevant if it is always repairing next the asset that will do the combat force the most good in terms of its weapon system availability goals; it is timely if it makes the repairs expeditiously, especially for those items most urgently needed; it is robust if its performance holds up in the face of uncertain repair demands.

Research Purpose

This report documents research conducted in Project AIR FORCE on contractor responsiveness, especially timeliness, in the repair of F-16 avionics components; it offers suggestions for policies that promise enhanced responsiveness. When an aircraft component is believed to be defective, it is typically removed from the aircraft and sent to an intermediate maintenance facility, usually located at the same base as the aircraft. If it cannot be repaired there, it is sent to the depot, i.e., the Air Logistics Center, for repair. For the items we examined, the bases typically send the unserviceable assets to the depot even if they are repaired contractually. The depot then sends them to the contractor and the contractor typically returns them after repair to the depot rather than directly to a base. The total elapsed time from the removal of a component from an aircraft until it is returned in serviceable condition to a base is commonly called *repair cycle time*. In this document, we are especially interested in the elapsed time from induction of an unserviceable asset by a contractor until its shipment as a serviceable from the contractor repair facility. We refer to this elapsed time as the *repair flow time*.

It includes the actual repair time of the item, which may be short, as well as all of the other processing times associated with its repair, including time awaiting parts. The principal difference between repair cycle time and repair flow time lies in the retrograde and order-and-ship times. The retrograde time is the time required to ship an asset from the base to the repair site. Order-and-ship time is the time from a base's requisition for an item to its receipt at the base.

Responsiveness is also a consequence of all the policies that govern contract repair, made by a number of directorates and organizations. As such, it is an important indicator of whether these guidelines and their execution work well together. A common barometer of performance used for contractual repair policies is cost. But with the substantial changes in the force structure occurring contemporaneously with the Air Force facing a much more dynamic and ill-defined threat than before, cost alone is not a sufficient guide. Moreover, a connection exists between repair flow times and pipeline costs. So, too, are there connections between the contract maintenance policies used for these assets and their implications for the overall support goal.

Using our observations in this work, we developed several suggestions as well as a set of hypotheses that need to be analyzed further and extended to all ALCs. As we will show, the current contractual repair management system has several important characteristics that inhibit its cost effectiveness.

Research Approach

We began this research to explore the feasibility of extending to contractors one concept developed for organic component repair processes to enhance their overall responsiveness. This concept, called DRIVE, operates on the principle that a facility faced with choices of what to repair next should give priority to those tasks that will contribute the most to weapon system availability over a specified planning horizon (Abell et al., 1992). Further, this concept proposes to distribute serviceable assets to bases in such a way as to maximize the probability of achieving specified aircraft availability goals on a worldwide basis.

The Air Force has adopted the concept of establishing repair and distribution priorities for some components repaired organically. The Air Force could extend this approach to components repaired contractually if it knew when serviceable assets emerged from contractual repair. As we delved deeper, however, we discovered that contractor responsiveness was a serious matter and deserved further attention. Indeed, addressing contractor responsiveness ranks higher in importance than anything else we can think of when it comes to private-sector sources of repair.

We focused attention on a small but important group of components for this research, namely, avionics components managed by the Ogden Air Logistics Center. Ogden has been a beta-site for many of the Air Force's concepts concerning enhanced depot-level responsiveness, particularly with regard to several initiatives developed by RAND. This history made the Utah-based ALC a natural choice for this work. Focusing on one ALC made practical sense. It made the problem more tractable and allowed us to examine one organization more carefully. This focus limits the direct transferability of our findings to other recoverables and ALCs. For example, the policy characteristics we note with avionics components we believe pertain more to those assets that fail unpredictably than to those that arrive at the depot on a scheduled, more predictable basis, such as engines. Nonetheless, we expect further analysis of other recoverables and ALCs to produce qualitatively comparable results.

Since many of the policies associated with contractual repair develop at the ALC level, some ALC-specific differences may also exist. However, the topics we discuss are of such a general nature that in all likelihood they apply to all the centers. Taking the approach we did has the unfortunate consequence of singling out Ogden from other centers. We believe it unwise and unfair to infer that the issues discussed pertain only to Ogden or only to avionics repairs. They are so fundamental in nature that logic argues otherwise, and to extend them to other commodities and ALCs is certainly worthwhile. Even though the observations we report here derive primarily from one ALC and one commodity, F-16 avionics, many of those we interviewed indicated that similar policies exist at the other Air Logistics Centers.

We used both qualitative and quantitative data. Anecdotal evidence came from a series of interviews that we conducted with individuals involved with contractors or contract repair policies, from the former Director of Contracting, to Item Management Specialists, who are internal customers of contractual repair, to several repair contractors who conduct a sizable volume of business with Ogden, to individuals at AFMC who provided valuable insight on how the government manages contractual sources of repair. We jointly analyzed two data sources for responsiveness information: the G009 and the G072D. We analyzed information from 19 repair contracts managed by Ogden. These contracts represent only a portion of the overall repair business between contractors and Ogden, but they were also the only agreements for which reliable data were available when we conducted this analysis. All 19 contracts reported activity on January 1, 1993, or later.

We were constrained in our analysis by data and data quality. No single information source exists to provide objective, accurate, and detailed data on

contractor responsiveness. Of the two systems used, neither was designed for the purposes for which we used it. Finally, although cost drivers implicitly surfaced during this analysis, quantifying them was beyond the scope of this research.

What Follows

In Section 2, we describe the repair process in general. Readers familiar with the steps of the process may wish to skip to Section 3, which sets out some of the important characteristics of the current system and suggests the need for changes in several functional areas associated with contractual repair. In Section 4 we discuss several policy issues affecting the responsiveness of contractual component repair. We offer several suggestions in Section 5 that we assert would improve the cost effectiveness of the processes and systems that support contractual component repair. Although we try to describe the key characteristics of contractual repair management, the reader is assumed to have some basic knowledge of the topic.

2. What Goes On with Contractual Repair?

This section describes the key processes involved in contractual repair, beginning with an estimation of the expected requirement over some planning horizon, usually one quarter, and ending with the receipt by a base of a particular serviceable. In later sections, we will discuss in more detail the implications of many of these processes or decision junctures and the policies that underlie them.

The decision to rely on a contractor for repair occurs early in the life of a weapon system. The Air Force determines whether it will develop organic repair capability or rely on contractor support throughout the aircraft's lifetime, although it cannot allocate more than 40 percent of its depot-level workload to contractors. In the event it decides to develop organic capability, the contractor may provide most, if not all, the depot-level support during the earliest years with interim contractor support (ICS). This approach of relying on contractor support during a time when technical changes in the new weapon system can occur ensures that maintenance procedures have settled down somewhat by the time the government invests in repair capability. The Air Force may also decide not to develop capacity beyond some level and may seek instead to buy the excess repair capacity needed to cover surges in demand from the private sector. Items that fall within that category are called split items (repaired by both the ALC and a contractor).

Contract Formulation

As a general rule, Ogden begins defining and drawing up a contract to buy repair services from the private sector months in advance of the need. The process of formulating the contract involves three steps: determining the repair requirement, identifying potential contractors, and negotiating the terms and conditions of the contract.

Determining Repair Requirements

The contracting directorate represents a critical interface between the government and the contractor. But the initiating action that leads to the need for contractual repair services originates with estimates made by the item manager on projected repair requirements over the planning horizon. For items

other than those related to avionics and managed by Ogden, the ALCs use the Management of Items Subject to Repair (MISTR) system to compute the estimated repair requirement (see Abell et al., 1992, for a more detailed description of the repair estimation process).

The requirements process begins with data contained in the Recoverable Consumption Item Requirements System (D041). This system computes requirements for recoverable spare parts and depot-level repair. It contains estimates of failure rates, repair times, and other item characteristics, and asset data from the depot and bases, although the asset position information is six to nine months old at the time the process occurs. Using these data, it projects repair requirements. Item managers (IMs) review these data and "scrub" the numbers by taking into account data or factors not known by the system, such as newly identified equipment problems that could change failure rates.

The IM reviews these data twice, once to identify inaccuracies and once to ensure that modifications were correctly entered; then the data are passed to the MISTR system (D073), which produces the X21 Report that accounts for assets due in from contractual repair sources. The X21 Report is used by the production management specialist (PMS) to finalize the quarterly repair requirement in coordination with the IM. The requirement that emerges from this process is called the *IM scrubbed requirement*.

Next, representatives from the organic repair facility and the PMS meet to help negotiate the allocation of workload. For split items, the PMS will allocate to the organic shop all the workload it can absorb over the planning horizon. He will then decide how much of the remaining repair work to buy from the contractor. If no organic capability exists and the item has only a contractual source of repair, then the scrubbed requirement represents an estimate of the services the ALC will seek to buy from the private sector. In the event the requirement exceeds financial resources, which is now often the case, then the center must develop priorities for what it will repair during the quarter.

Because estimates of repair requirements have such an important effect on the type of contract negotiated, it is important to summarize the most relevant features of the processes described above:

- The traditional estimation processes used by the Air Force do not account explicitly for aircraft availability.
- An information system that bases its recommendations on considerations of worldwide aircraft availability is currently used in a limited way, but not for items repaired by contractors. Section 4 takes up this topic in more detail.

- Two important sources of uncertainty exist: demands that generate over the quarter, and the portion of the expected repair requirement actually bought. These uncertainties help shape the type of contract the ALC ultimately negotiates with the contractor.
- For those items requiring a contractual source of repair, the PMS will ask the contracting directorate to negotiate a particular type of contract for a list of national stock numbers (NSNs) and expected quantities. The greater the uncertainty associated with the expected quantities, the more likely the ALC will seek highly flexible terms that trade off unit repair costs against up-front commitments on the eventual requirement. That is, if the quantities are highly uncertain, the ALC may try to fix the unit repair price based on some minimum quantity while avoiding any guarantee on the overall repair requirement.

Potential Contractors

A contractor must meet a number of qualifications for consideration. He must be certifiably capable of performing the repair up to the quality standard required. If the potential contractor is not the original equipment manufacturer (OEM), then the Air Force can provide him with technical data to develop repair procedures and plans. The Air Force can buy technical data from the OEM, usually at the end of an acquisition program, for its own engineering purposes. However, it sometimes forgoes purchasing that information if the original procurement program runs into budgetary problems. In that case, if a new contractor wants to qualify for the repair work, he may try to reverse engineer the component to infer an appropriate repair process.

Past policies, now subject to change, called for stepped-up competition among sources of repair based on the principle that greater competition would lead to reduced repair costs. In response, Ogden created an office of competition advocacy, separate from contracting, that would identify new contractual sources of repair. This office originally set up a buffer between those who had to find new sources of repair and those who would award and manage contracts. Within the last several years, the center also established a business and marketing office designed to help bring new repair workloads to the depot through participation in public/private and public/public workload competitions. Interest in conducting competitions between the government and the private sector has waned, indicated by a DoD directive to put further

public/private competitions on hold.¹ But private/private competition for those workloads requiring a contractual source of repair still continues. Since the decline in defense procurement, former suppliers to original equipment manufacturers have tried to win more of the repair work. Just as new sources of repair are coming on-line, others are going out of business or turning to other markets. Thus, some level of flux exists.

In some cases, the government opts to contract with the OEM or one of the subcontractors to act as a clearinghouse for vendors they have partnerships with. The ALC can reduce the number of contracts it manages by following this practice. On the other hand, the OEMs note with chagrin that the ALC will contract directly with one of its vendors if the government believes it can get a better deal by doing so. Heightened competition has led to lower repair prices in some cases along with new sources of repair.

Contract Negotiations

Government contracts must adhere to stipulations set forth in the Federal Acquisition Regulations (FARs). The FARs define boundaries, conditions, and constraints a government entity must abide by when contracting with the private sector for goods and services. They have far-reaching effects. In general, the legal constraints they impose limit the decisionmaker's flexibility, but the interpretation of these constraints can change.

ALCs may negotiate at least five distinctly different kinds of contracts, with many variations within each category. They are:

- Fixed-price: quantities, prices, and schedules are fixed, such as for off-the-shelf, commercially available items.
- Indefinite-delivery: at least one of the major parameters—quantity, price, or schedule—is uncertain; the requirements contract falls under this heading.
- Incentive: the contractor is rewarded for improved performance; used in cases where cost and performance data are not sufficiently known.
- Time-and-materials or labor-hour: the government pays according to rates rather than a flat fee, usually with the condition of a price cap.

¹Communications with Ogden, November 1994. Also, see Camm et al. (1994) for further discussions on past public/private competitions.

- Cost-reimbursement: the government pays as it goes; this represents the most uncertainty in cost and performance, such as in research and development (R&D).

Important considerations going into a negotiation concern the availability of cost data, estimation of performance or cost risks, the number of qualified companies bidding on the workload, and the certainty of the workload requirement at the time of the negotiations. If most of these elements are known beforehand, then the contract type will fall closer to a fixed-price type. The number of competitors and the relatively low technical risk involved with such requirements allow the ALC to lock in prices. If, on the other hand, most of these elements are not known from the beginning, then the government may need to assume more of the risk and uncertainty, reflected in the type of contract used as well as the compensation awarded.

A type of indefinite-delivery arrangement preferred by Ogden is a requirements contract, which fixes the unit price but leaves the delivery (or repair quantities and schedule) uncertain. The government may use an incentive contract if it needs improved delivery or technical performance from a contractor and insufficient information exists on the technical or cost risks at the time of the award. By linking potential profits to performance, the contractor may be more strongly motivated to search for innovative solutions to particular problems. Time-and-material and labor-hour contracts offer useful vehicles in cases where the costs and requirements are unknown at the time of the award. Given such unknowns, the government agrees to pay as it goes. Finally, cost-reimbursement contracts pertain primarily to R&D facilities.

The entire negotiating process, from beginning to end, can easily span 10 months. If the government has no previous experience with the contractor, it can take longer. Contracting officers find themselves in a virtually constant state of discussion if the contract runs for only one year. To ease this burden, particularly with contractors that have a long-standing relationship with the ALC, a multiple-year contract may be negotiated, with an option to renew at the end of the annual period. This approach benefits both the contractor and the ALC. If the contractor works with the ALC year after year, the two parties may agree to review unit repair prices before the renewal period with the understanding that if cost estimates are off for one year, the opportunity to shift them up or down is available.

Increased competition has altered these arrangements somewhat, making long-standing relationships less certain. Paradoxically, competitions for workloads

between the public and private sectors in recent years increase the contractor's risk in sharing cost information broadly with the government.

Source Selection

The government will select a source of repair from among those contractors who have proved they are qualified to do the work. Usually it will award the work to only one contractor. In the past, the award process involved straightforward choices: low bids had a distinct competitive advantage over all others, with low bids more often than not walking away with the prize. Two facts may explain this phenomenon: repair costs are easily measured and awards based on low repair costs are easily defended, if necessary. Unfortunately, unit repair costs are not necessarily the best indicators of low overall costs.

With the recognition that quality and responsiveness also matter, the government is trying to broaden its selection criteria. Importantly, the Office of the Secretary of Defense (OSD) has issued policy directives that will require the services to consider past performance. AFMC is evaluating several information systems to aid in this effort. (See Appendix A.)

Contract awards occur year-round. They may begin smoothly when the old contract expires. If they do not, the government may choose to extend an older agreement to cover the gap in service that might otherwise exist.

Workload Formulation

The Air Force conducts four repair cycles annually. Thus, on a quarterly basis, the ALC projects the repairs likely to be generated over a three-month period and the repair backlog, to come up with a repair requirement. It allocates the workload associated with that requirement to organic and contractual repair activities, and for those items repaired on contract, it executes delivery orders.

Determining Quarterly Repair Requirements

The IM's final estimate, accounting for all that he or she knows about the component itself, goes to the PMS. Once the IM and the PMS have agreed jointly on the list of NSNs and quantities to be repaired, workload negotiations and allocations take place.

For many components, DRIVE develops a list of NSNs, quantities, and priorities, based on maximizing the probability that aircraft will meet specified aircraft

availability goals in peacetime and wartime worldwide. This system is described in detail elsewhere (Abell et al., 1992).

Workload Allocation

The PMS and the organic shop chiefs negotiate workloads for split items, giving the organic shop workload up to its capacity and contracting out any remaining amount. Whether the ALC buys the entire remaining amount depends in part on whether it has sufficient funding to cover it. Very often the government must decide to repair only a subset of items during the quarter. The unmet requirement carries over to the next quarter's repair computation. If the item has only a contractual source of repair, some further refinement of the requirement may occur, particularly if budgetary constraints exist, as are now common. One form of setting priorities for contractual repairs is to develop "critical item lists" of those items deemed to contribute to mission capability (MICAP) conditions or to be in chronic short supply relative to demand. The ALC would fund those repair requirements first.

The reason given for the allocation policy for split items—to give the base workload to the organic source of repair and any excess to the contractor, subject to funding availability—centers on a measure of performance used by the ALC: labor effectiveness rates. The conventional wisdom throughout AFMC argues that organic shops perform repairs cheaper and therefore should be preferred. Moreover, the government has an interest in keeping its resources fully employed. For both these reasons, the tendency to allocate workload for split items to the organic source of repair is very strong.

Executing a Delivery Order for Component Repair

The delivery order technically represents an amendment to the original contract. Essentially, it authorizes the contractor to conduct a repair upon receipt of the unserviceable. These orders are usually a page or two and list a set of NSNs, along with required quantities. They may also specify funds for contractor acquired property (CAP) or for reporting requirements.

Because the order represents a legal document, complete with signatures, the contractor will usually wait until he has both the delivery order and the unserviceable asset. Although the ALC tries to coordinate these two events, one may arrive before the other.

Repair Process

Four major events in the repair process occur at a contractor's facility: unserviceable arrival, component induction into repair, the repair itself, and shipment of the serviceable. We describe each of these below.

Unserviceable Arrival at the Contractor

Unserviceables arrive at the contractor from one of three sources: the ALC, an Air Force base, or a contractor who must transfer unserviceables stored at his warehouse to another contractor, e.g., during the start-up of a new contract. In most of the cases we observed, the unserviceable came from Ogden. A base will sometimes ship an unserviceable directly to the contractor in the event of a peacetime emergency, such as an aircraft that is not mission capable because it lacks a mission essential component. Those we interviewed told us that unserviceables are batched. The depot sometimes opts to wait until a number of assets accrue for a particular contractor before shipping them. The goal in this case is to minimize processing and transportation costs. Batching may also come about because either funding constraints or the depot's quarterly workload negotiations lead to repair decisions made simultaneously across groups of assets.

Once the asset arrives at the facility, it is logged in. If the contractor has both the delivery order and the unserviceable, the clock starts. Most of the contracts we observed required contractors to deliver serviceables within a mean time of 30 or 45 days, barring unavailable repair parts, discussed later.

Induction into Repair

Two types of repair lines exist: the production line itself, if the contractor is an OEM, or a dedicated repair facility. Because repair is generally less profitable for the contractor than manufacturing, given the choice, he will prefer to schedule repairs around production. If the contractor has a dedicated repair facility, inductions need not compete unduly with other jobs. However, such a facility implicitly depends more on a base workload to meet its costs in each period.

Contractors prefer scheduling repair inductions to produce smooth workloads. If workloads vary greatly from one period to the next, the contractor must find other tasks for his workforce during periods of low demand, and if demand stays low he may be forced to lay off some workers, or during high demand periods he may need to hire and train new workers.

Repair

The contractual clock counts the time it takes for the contractor to perform all necessary actions. If, in the course of a repair, the contractor requires a repair part that is owned and managed by the government but not available on the contractor's shelf, the contractor orders the part and the clock stops. Once the part is received, the clock starts again. The elapsed time between these two events is time awaiting parts (AWP). AWP time represents unproductive time for the asset and both organic and contractual sources of repair would prefer to minimize it.

The repair process consists of three primary tasks: diagnosing the fault, repairing the fault, and retesting the serviceable. Avionics components are usually repaired using automated test stands that run through a series of diagnostics, or by a highly skilled technician who can perform similar tests. The component, called a line-replaceable unit (LRU), typically contains one or more shop-replaceable units (SRUs). Often an LRU fails because at least one of its SRUs has a problem. Likewise, a shop may repair an SRU by replacing faulty sub-SRUs or component parts. At its core, the repair process consists of diagnosing faults and replacing unserviceable components with serviceable ones.

When the need for, say, a serviceable SRU arises and none is available, the LRU must be removed from the test stand or work station and stored until an SRU is available. When the serviceable SRU is in hand, the LRU must be reinstalled on the test stand and diagnosis and repair started over again. SRUs are repaired by replacing or repairing sub-SRUs.

Ideally, the Air Force would allow a contractor to stock a small inventory of repair parts and serviceable SRUs to draw from while conducting an LRU repair. As these parts and serviceable SRUs were used, the contractor would replenish this stock with new supplies or by the repair of the removed SRU. By keeping these parts and rotatable spares on hand, the repair process is not interrupted. In reality, however, Ogden says it often finds itself short of funds and unable to lay in adequate material support.²

Some contractors repair SRUs; others do not. If rotatable spares are not available and the contractor is authorized to repair faulty SRUs, these SRUs will be repaired contemporaneously with the LRU. Any delays in bringing the SRU to serviceable condition lead to a delay in the LRU repair. To make matters more complicated, in the repair of avionics components, the software diagnostics

²Communications with Ogden ALC, November 1994.

depend on a particular sequence of successful tests of SRUs. Thus, compound delays may occur in a single LRU repair.

As mentioned previously, the contractor usually has on hand some material support in the form of repair parts or consumable items. If the parts needed for a repair are not on hand, they must be ordered. The government purchases and manages many parts, commonly referred to as government furnished material (GFM), and is often the source of supply. When that occurs, a Military Standard Requisitioning and Issue Procedures (MILSTRIP) action takes place. In the event that the government discovers it cannot fulfill the request and additional supplies are not expected for some time, it gives the contractor permission to purchase the material on its own with Air Force funds. This material is then called contractor acquired property (CAP). The contractor may purchase the material from himself, if he is the manufacturer, or from another vendor.

In the past, the government specified that the repair parts purchased for one aircraft program could not be shared by another. Thus, it is possible, and we were informed it has actually happened, that a contractor might have to requisition from a government source repair parts that are in stock in his own facility but purchased for other aircraft programs. In this case, of course, the parts are common to more than one aircraft. The rationale for sequestering these parts is to give priority to some weapon systems (such as the B-1 over the F-16).

Finally, contractors may vend out part or all of the repairs, as allowed by the contract. In this instance, the government does not deal directly with the vendor. Vendors quite often face similar constraints as contractors, in wanting smooth workloads, etc. Usually in this type of arrangement, contracts allow for longer negotiated repair flow times.

Shipment from the Contractor

Ogden requires that most contractors ship all serviceables to the depot upon completion of the repairs, except those applied against a MICAP. Some contractors prepare the shipping instructions and paperwork early in the repair process. If Ogden requests that the contractor ship the serviceable directly to a base, the government will issue amended shipping instructions that give the contractor the legal authority to respond to the immediate need and waive the usual shipping directions set by the contract. Usually these instructions are used only for MICAPs. If the contractor uses a vendor to conduct the repairs, the same information is passed to the vendor.

3. How Responsive Is Contractual Repair?

The customer looks for at least three things in depot-level support: (1) getting the service (2) at an affordable price (3) when needed. We made the argument that the Air Force's cost of support is a function of repair cycle time, among other things, because repair cycle time affects the number of items in resupply; i.e., the Air Force must invest in a larger inventory of items to fill the pipeline in cases where repair cycle times are very long. The third aspect noted above, getting service when needed, speaks to the depot's ability to respond quickly to urgent, unanticipated demands. Long repair times limit that flexibility.

Coronet Deuce, a series of three demonstrations at Ogden of the concept of two levels of repair for avionics components, provided evidence that when greater emphasis is given to timely repairs, organic shops can reduce their repair flow times very dramatically. Coronet Deuce led the Air Force to conclude that the deletion of F-16 intermediate-level maintenance capability would not jeopardize aircraft availability. By making a few key changes in the repair process, the Ogden F-16 LRU repair shop is now able to move over 75 percent of the repaired assets through depot repair (from induction to shipment) in 7 days or less, and has achieved a mean repair flow time of about 12.5 days including AWP time. Such improvements in organic operations raise the expectation that the Air Force will move toward similar initiatives in contractual repair operations.

In the previous section we noted that contracts usually specify a repair flow time requirement that the ALC expects the contractor to meet, such as a mean time of 30 or perhaps 45 calendar days. In this section we take a closer look at some of Ogden's contracts starting in January 1993.¹

If we consider one aspect of repair responsiveness, i.e., its timeliness, our research indicates that contractual repair flow times are long. Moreover, an inspection of the contracts and information gained through interviews indicates that the Air Force does not emphasize responsiveness. We surmise this because of the absence of language pertaining to responsiveness in the contracts, apart from the mean repair flow time specified in each contract. Furthermore, in trying to quantify the repair flow times being achieved in current contracts, we found no complete data source that yielded this information. The Air Force must

¹For details on our data and analytic approach, see Appendix C.

emphasize responsiveness with its contractors if it hopes to significantly reduce repair flow times as it did in its organic operations.

In this section we discuss some of the issues that can contribute to long repair flow times. Our discussions include an overview of repair flow times, ALC repair requirements estimation processes, how split item repairs are allocated to their two sources of repair, and legal constraints.

Repair Flow Times

In the nineteen contracts we analyzed, we found long repair flow times that on average exceeded the flow time requirement set by the contracts.² Repair flow times, remember, refer to the number of calendar days it takes for a contractor to ship the serviceable asset from the time he inducts it for repair. Usually the contractual language used allows contractors to exclude AWP time when reporting repair flow times. Thus, if an item with a 30-day negotiated flow time requirement stays at a contractor's facility for 40 calendar days, but spends 10 days of that AWP, then the ALC would consider this item to have met the 30-day requirement. If, however, the item spent 40 days at the facility and no time in AWP, then the contractor would need to repair the remaining items in sufficiently short times to stay within the 30-day mean repair time.

The above example illustrates one problem we encountered in analyzing the data sources recommended to us for this research: Given only the dates associated with asset receipt, induction, production, and shipment, we know nothing about the time assets may have been AWP. Contractors often supply the PMS with their own performance statistics, but the repair flow time they report is often clock time that excludes AWP time. According to Ogden, current reporting systems do not allow it to measure AWP time. Contractors report GFM status by hard-copy forms once a month or by an automated system described in Appendix C, G009.³ These data do not yield AWP time.

AWP problems often require the involvement of the ALC to resolve. The ALC will get involved if a MICAP occurs and the lack of parts that are on order prevent the contractor from completing a needed repair. The contractor may also ask for help from the PMS to get GFM released if a shortfall in repair parts prevents him from meeting an internally imposed schedule, such as achieving a

²We discuss reasons for focusing our analysis on this subset in the next section. In essence, we were restricted to these items because of data quality problems.

³Air Force Form 412, entitled the Government Furnished Material Report, is used by some contractors to report their GFM inventories. The G009, used typically with large repair contracts, produces an automated version of this form.

certain revenue stream. Although teams at Ogden are beginning to monitor AWP problems for some organic repair operations, discussions continue on how this might be done for contractual repair.

During this research we were unable to identify the time an asset was AWP. The value of coupling repairable AWP time with repair flow time is enormous: It is one of the few ways to track contractor performance in an objective, quantitative manner. More importantly, it provides the Air Force the information it needs to help contractors reduce AWP and repair flow time.

From the data sources available to us, we could only retrace what happened, not why. Accordingly, we could not independently verify contractor performance or identify the root causes that led to long flow times. Nonetheless, we believe that the calendar time these assets spent in repair at a contractual facility stands on its own as a legitimate measure of how the system behaves. After all, from the customer's point of view, the real time it takes to get service is what matters most.

For the weapon system we examined most closely, the F-16, the Air Force bought between 37 and 56 percent of its depot-level support from contractors during the 1985–1988 period.⁴ We found that, of those assets repaired under the terms of our 19-contract sample with reporting dates that started in January 1993, about half spent more than 60 days in some contractor's hands (see Figure 3.1). The data, shown as cumulative percentages of total assets, indicate calendar days from the time an unserviceable arrives at a repair facility to when the contractor ships the asset in serviceable condition. A fifth of the recoverables repaired under these contracts spent more than 100 days at a contractor's facility.

As already noted, the three principal stages involved in contractual repair are waiting for induction into repair, the repair process, and waiting for shipment from the contractor's facility. We separated the repair data into four categories: (1) F-16 NSNs with 45-day negotiated flow times (NFT); (2) F-16 NSNs with 30-day NFTs; (3) a combined set of the first two groups; and (4) the total sample of items. The total sample included some non-F-16 NSNs.

Of the three stages described above, the assets we analyzed spent the bulk of their time in the repair process while at a contractor's facility (see Figure 3.2). For this group of contracts, recoverables with 45-day NFTs spent a mean time of 77 days at a contractor's facility. These same assets spent an average of 71 days in

⁴Source: Weapon System Cost Retrieval System (WSCRS). Karl Hoffmayer, a RAND analyst, kindly supplied these cost data. Unfortunately, WSCRS data from 1989 to the present were unavailable at the time of this analysis. Inflation factors were provided by OSD.




Figure 3.1—Repair Flow Times for 19 Contracts with Initial Reporting Beginning in January 1993

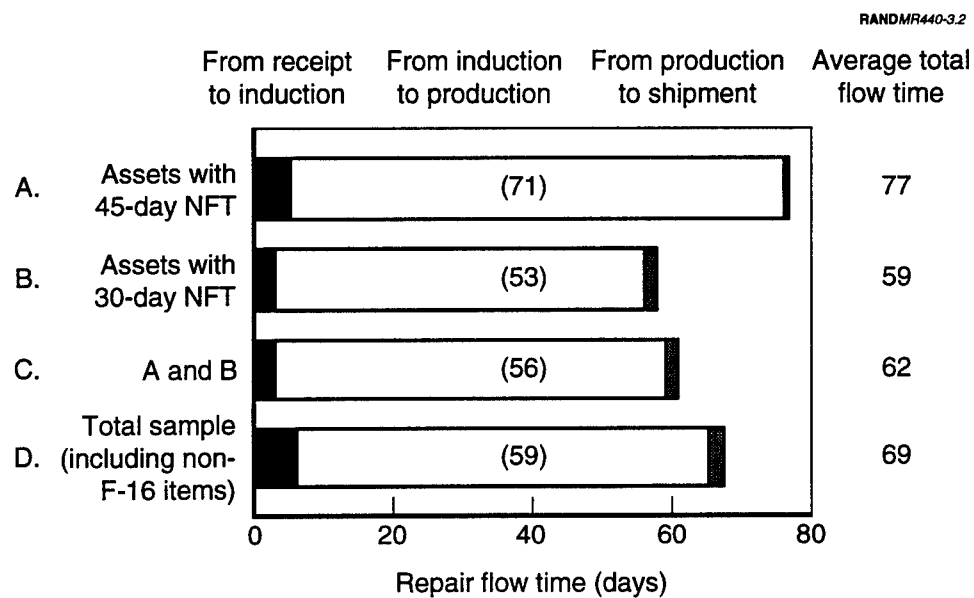


Figure 3.2—Mean Repair Flow Times for 19 Contracts with Initial Reporting Beginning in January 1993

the repair process, measured from when they were put into repair (induction) to when they emerged as serviceable (production). Another class of assets with 30-day NFTs spent a mean time of 59 days in the contractor's hands and of that, 53 days in the repair process.

The contracts we examined held the contractor accountable from receipt of the delivery order or the unserviceable asset, whichever occurred later.⁵ As mentioned earlier, a delivery order represents the legal authority to begin work on a prescribed list and quantity of assets. Assets may arrive before the delivery order does, as can easily happen when a base ships an unserviceable directly to the contractor. Other events can delay the contractual clock from starting, as when large batches of assets and orders arrive together. Batching requirements can make it difficult for the contractor to conduct repairs cost-effectively. An induction rate that meters assets into the repair process can smooth workloads and is preferred by contractors, particularly if they operate a dedicated repair facility. Delaying induction and the contractual clock requires Air Force approval, however.

We speculate that long repair flow times derive largely from AWP problems or the lack of serviceable SRUs and other repair parts (see Figures 3.3 and 3.4). It is also possible that the contractors themselves bear some of the responsibility for these flow times. In any event, long repair flow times imply high inventory costs to fill repair pipelines. Given the long flow times we observed in our contract sample, the inventory implications should give the Air Force an incentive to think about how it manages contractual repair, particularly if it wants to allocate more and more of its workload to the private sector.

Interestingly enough, in both the 30-NFT and 45-NFT cases, we found no correlation between repair costs and flow times. If one assumes that items with higher repair costs are also those that cost more to buy new, then the absence of any correlation indicates that contractors treated expensive items no differently from inexpensive ones. One would like to see expensive items managed more intensively (i.e., better material support and less AWP time) to keep inventory investment costs low. We did note that, among items with a 45-NFT requirement, those with higher costs for repair parts also tended to have longer repair flow times.

⁵We began our measurements with the time of induction rather than unserviceable arrivals to account for legitimate delays deemed to be within the control of the ALC and not the contractor (delays in getting delivery orders, negotiated induction rates, etc.). As noted earlier, the cycle time between production and shipment is small—about 2 days on average for our total sample. The clock does not stop until the contractor finally ships the serviceable.




Figure 3.3—Repair Flow Times for Items with 30-Day Negotiated Flow Times




Figure 3.4—Repair Flow Times for Items with 45-Day Negotiated Flow Times

Mean Repair Flow Times Do Not Tell the Full Story

Using mean flow time as a standard is far less demanding than requiring that flow times for all items meet an absolute standard. We can use the data in Figures 3.3 and 3.4 to show how big the difference can be. Only 30 percent of the

items in the sample in Figure 3.3 could have met an absolute standard that all items have repair flow times of 30 days or less; almost 60 percent of these items could have met the standard that these items, taken together, have an average repair flow time of 30 days. Using a standard based on an average effectively doubled the absolute time standard that items could meet, from 30 to 60 days, and still comply with the standard.

For the items in Figure 3.4, the difference is even larger. Almost 40 percent of the sample could meet an absolute standard of 45 days. Using an average standard of 45 days is equivalent to extending the absolute standard from 45 to 110 days.

Important Implications of Uncertainty in Repair Requirements

As we pointed out before, component repair contractors are vulnerable to uncertainty in repair requirements. In the context of this discussion, we use the term requirements to mean the number that finally gets written into a delivery order, the quantity the contractor will actually try to repair given that he has or receives at least that many unserviceables. The uncertainty that pervades the repair contracting scenario before the signed delivery order being given to the contractor derives from several sources: (a) uncertainty in unserviceable generations, (b) uncertainty in repair requirements determination from quarter to quarter, (c) uncertainty in funding from quarter to quarter, and (d) uncertainty in repair parts availability. If, at the start of a contractual period of a year, say, the actual number of unserviceables that would be delivered to the contractor each month or even each quarter were known, and the funding of all those repairs were somehow guaranteed, the contractor would probably be willing to negotiate a very different set of contractual arrangements than he would with the uncertainties that now pervade the contracting environment. On the other side, the Air Force would not need to pay for the uncertainty against which the contractor needs to hedge in negotiating the terms of the contract.

What the contractor fears is very real. We were told that there have been cases where the Air Force ran out of money before the end of a fiscal year and had to tell some of its contractors to stop work, and that there are many examples of cases where the Air Force specifies planning quantities for repair contracts that far exceed the total requirements specified in later delivery orders, eventuating in the need for fewer resources than the contractor had originally planned for.

Given such circumstances, it is little wonder that the Air Force pays a premium price for contractual repair. In its use of minimum repair quantities in requirements type contracts, the Air Force virtually assures a high unit price for

component repair. Yet, the contractor must insist on some protection against such unpredictable and damaging behavior. This unhappy situation derives largely from two factors: (a) the uncertainties already mentioned, and (b) the Air Force's practice of allocating workload to its organic facilities up to their capacity and the remainder to contractors, which, as we pointed out earlier, allocates the more uncertain workload to the contractor.

It needs to be understood better what price the Air Force pays for uncertainty in component repair contracting. Our hypothesis is that the organic repair facility is better able to cope with uncertainty in repair workload than contractors typically are. One important reason for this is that organic repair facilities usually have greater scope of repair than contractors do. By that we mean that the organic facility is able to repair a larger variety of component types and is therefore less vulnerable to variability in workload than are facilities with less scope of repair. We were not able to evaluate this hypothesis, but we judge it to be an important issue in reengineering the management of contractual repair.

The only certain thing in the contracting scenario is uncertainty, and the Air Force needs to find ways to reduce it. Although uncertainty in unserviceable generations will always be a factor in this problem, funding uncertainties need not be.

Perceptions of Contractual Repair Responsiveness

We found that within the Air Force community, among contractors, Ogden, and AFMC, contractual repair responsiveness is viewed as an important issue. Indeed, one person volunteered to investigate those NSNs we identified as having the longest repair flow times.

Perhaps not surprisingly, contractors say they view themselves as responsive, particularly when it comes to MICAPs. Although repair contracts contain language that requires contractors to give top priority to MICAPs, coordination between the ALC and the contractor must still take place, often conducted by phone or fax. Thus, they see themselves as problem-solvers and are willing to go the extra mile to provide their customer, the government, a good product.

Contractors quite legitimately measure themselves against the terms of their respective contracts. Many of them do not have adequate rotatable spares on hand (also known as serviceable SRUs) or stocks of consumables. Some contractors say they have few incentives to push for beefed up material stockage policies as long as their contracts allow the clock to stop when an asset became AWP.

The ALC relies principally on its staff to monitor contractor responsiveness, as defined by contracts. Some contractors report elapsed repair flow times to automated data systems; others report this information via hard copy worksheets. The dollar value of the contract is one factor the ALC uses to decide the appropriate reporting medium, i.e., not all contracts require G009 reporting. Private-sector repair facilities report GFM status via hard copy forms or the automated system, depending on the particular contract. AWP status is more difficult to infer. Shipments of serviceables from the contractor are reported to the IM. Importantly, these actions initiate payments.

For the most part, the government requires contractors to report transactions on the basis of NSN instead of NSN and serial number. Engines are an exception to this rule (although an engine that has just had a major overhaul may have only its identification plate remaining from its original set of components). The argument against serial-number reporting for most other components rests, in part, with the implied data storage requirements associated with detailed information and the data entry effort involved. Although data reporting on an asset-by-asset basis would increase data entry and storage requirements, our judgment on this matter is simply that the data are worth more for monitoring the contract repair process than they cost in terms of these requirements.

A consequence of reporting repair processing events to several different directorates within the ALC on the basis of NSNs only is to make it extremely difficult to monitor contractor responsiveness in an objective, accurate way.

We next discuss additional implications of our observations about the current system and explore some solution directions.

4. Policy Directions for Enhancing the Responsiveness of Contractual Repair

In the previous section, we discussed current contractual repair responsiveness and showed that contractors for the F-16 currently experience far longer repair flow times for depot-level repair than the Coronet Deuce demonstrations have shown are achievable. To help the Air Force reduce contractual repair flow times and enhance the responsiveness of contractual repair, in this section we suggest some policy directions that may be helpful. Our suggestions indicate directions of change, not necessarily specific solutions.

Improve Organizational Arrangements for Contractual Repair Management

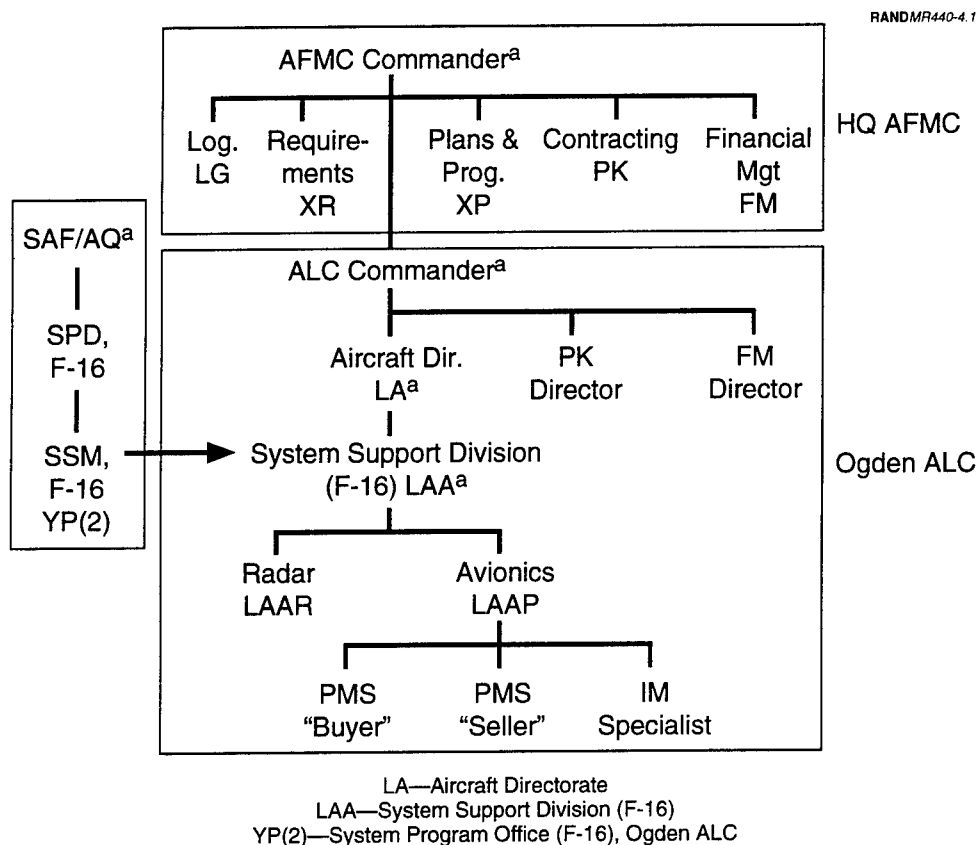
The first step toward new policy is to define an institutional advocate and sponsor for needed change. Who should take responsibility for developing policies that can improve the responsiveness of contractors? Many instinctively point to officials responsible for contract policy. They, after all, write the contracts that set the terms that contractors work against. But contracting officials do not set substantive policy. Rather, they seek to implement substantive policy set elsewhere. They ensure that any contracting policy satisfies legal and regulatory constraints and places the Air Force in the best position, relative to a contractor, that these constraints allow. In this sense, contracting officials play a role similar to that of a lawyer seeking the best legal way to realize his client's intent. The client must still retain responsibility for stating that intent.¹

In this setting, the relevant "client" is presumably the office responsible for buying services from contractors. Several years ago, that office would naturally have been Material Management (MM) within the Air Force Logistics Command, which bought repair services from Maintenance (MA). The DCS/Logistics (LG) at Headquarters, AFMC, has inherited the functions formerly performed by both of these organizations. But it primarily represents the organic provision of

¹Financial management (FM) similarly sees itself as a service organization with little or no authority to change substantive policy by itself.

depot-level services. Also, two recent systemwide changes have complicated any effort to find a single advocate for the buyer of depot-level services.²

First, integrated weapon system management (IWSM) has refocused depot-level services around individual weapon systems (AFMC, 1993b). System support managers (SSMs), responsible for support policy on individual weapon systems, report through their system program directors (SPDs) to the Office of the Secretary of the Air Force (SAF), Acquisition. (Figure 4.1 illustrates these relationships for the F-16.) No single office in an ALC or at Headquarters, AFMC, can speak for these individual offices responsible for developing contractual repair policy.



^aNot all subordinate offices are shown.

Figure 4.1—Organizational Chart of Key Contract Maintenance Decisionmakers

²While the Air Force was pursuing public/private competitions during 1991–1994, OSD policy required that it separate organic buyer and seller organizations to ensure that organic sellers did not gain an unfair advantage over contract sellers of the same services. This requirement further complicated LG's role as a single advocate for buyers when its principal function was the provision of services.

Second, even without this change, the inclusion of depot-level recoverables (DLRs) in the Defense Business Operations Fund (DBOF) has linked funding for depot-level repair from AFMC with DLR funding for the major operating commands. Each of these commands has at least a portion of the financial role typically associated with a buyer. And each is individually postured to play a potentially more active role in setting the terms under which depot-level services are provided.

To the extent that any single substantive policy office within AFMC might take the lead in promoting greater responsiveness from contractors, the most likely candidate appears to be the XR office at Headquarters, AFMC. XR has indicated that today its designated responsibility for contract repair is not clear.³

Until a clear proponent for improved responsiveness is established, the contracting community is likely to dominate any policy discussion by default. Contracts will continue to be written, and the community's strong preference for proven contractual forms and language will discourage change unless someone presents a compelling case for change. This perspective will make it especially difficult for individual policy entrepreneurs in various parts of the Air Force to effect significant change. The Air Force is most likely to break its traditional approach to contracting if it makes a concerted and focused effort to do so. Contracting officials will play a central role in such an effort to ensure that contemplated changes are legally feasible. But someone else must provide the motivation for policy change and the energy to ensure that it occurs.

Improve the Responsiveness of Contract Repair

Perhaps the most compelling single observation in this work is that the Air Force needs to improve the responsiveness of its repair contractors. There is persuasive evidence from the Coronet Deuce experience that organic repair shops can improve their responsiveness dramatically and, equally important, can do so at very little cost. While Ogden is writing repair contracts for repair flow times of 30 or 45 days for F-16 avionics LRUs, its organic repair facility is repairing the same kinds of items and achieving repair flow times (from induction to shipment, including AWP time) of 5 days or less 68.4 percent of the time. Moreover, it is achieving an average flow time of 12.5 days in contrast to the 59 days and 77 days achieved by contractors on 30-day and 45-day NFT items, respectively.⁴

³Communications with AFMC, January 1994.

⁴Statistics extracted from the Material Tracking System (MTS) were provided by Ogden Air Logistics Center. The MTS tracks by serial number organic depot repair flow time, including AWP

As we pointed out before, when it comes to spares investments, time is money. The shorter the depot repair pipeline, the lower the spares capitalization cost. Moreover, desired levels of relevance and robustness in depot repair are much more difficult to achieve in the face of long repair flow times.

Reduce AWP Times at Contractual Repair Facilities

There are two issues involved in reducing AWP times in contractual repair. The first involves measuring times accurately for LRUs and SRUs. The second involves the development of a more effective stockage policy for repair parts to support contractual repair. As we pointed out in Section 2, lack of a serviceable repair part when needed not only causes the end item to wait, it also induces inefficiencies in the repair process, a compelling argument in favor of initiatives to reduce the duration and frequency of AWP incidents.

Although policies exist to lay in repair parts and material on some contracts to some degree, the ALC does not manage material support extensively except in cases of more predictable demands (such as scheduled engine repair).⁵ In general, the government does not stock adequate supplies of rotatable spares or serviceable SRUs at the contractor's repair facility, nor do contractors have much discretion over buying parts in advance of actual need, except for some inexpensive parts, such as nuts, bolts, washers, etc. But even if the contractors were given more discretion, they currently do not have many incentives to avoid AWP occurrences, since AWP time is not counted in performance measures.

To appreciate the benefits of laying in parts and rotatable spares, the Air Force needs to first measure AWP times to understand better the cost-effectiveness of any proposed material support policy. Emphasizing responsiveness and measuring all segments of the repair pipeline would motivate the rethinking of material support policy for contractual repair. A sensible material support policy should involve the contractors themselves to an important extent.

To the extent legislation prevents the Air Force from addressing this critical issue, the analysis conducted to demonstrate the benefit of tailoring material support more closely to the nature of the repair process could also be used to argue for lifting some of the most costly constraints. Establishing the elapsed times

time, for each asset in the database. Ogden developed the system originally in support of Coronet Deuce. It is used today for avionics two-level maintenance recoverables, and for selected assets in the Air Force demonstration of Lean Logistics, now under way. Information was provided by communications with Ogden ALC on October 25, 1994.

⁵Items whose demands are fairly predictable, such as engines that are removed and overhauled after a certain number of flying hours, reportedly have much better material support policies. The best indicator of material support problems is flow times that exceed their contractual goals.

attributable to AWP conditions provides information critical to objective performance evaluation, but more importantly, to the improvement of contractor responsiveness. Unfortunately, the G009 system does not explicitly report detailed information on AWP events or their duration. We understand that AWP information does not link back directly with any data source that reflects information about contractual repair responsiveness. We suggest that the command consider building a crosswalk between the G009, which potentially provides a good source of contractor responsiveness data (Air Force Materiel Command, 1993a) and a reliable data source for AWP information. Further, this responsiveness information could be shared with the past-performance data system the Air Force is contemplating for repair contractors.

If responsiveness were an important concern of contractual repair management, it would probably motivate policy changes in such areas as material support and data reporting requirements. AFMC and the ALCs cannot responsibly enact new policies in an information vacuum, especially with a downsized fighting force that depends on cost-effective logistics support. For the Air Force to emphasize responsiveness throughout AFMC and the ALCs, it should begin to measure it explicitly.

Ship Assets Directly Between Bases and Contractors

Perhaps in the past one could argue that poor in-transit visibility made it necessary for all assets to pass through the depot on the way to a repair activity and back again just to take a physical count of what was in the system, but current information technologies make it possible to know where assets are without having to touch them. A direct payoff in passing data rather than the components themselves is enhanced depot-level responsiveness, lower inventory costs, and better service. Thus, the core responsiveness issues pertaining to distribution policies for serviceables and unserviceables are how to cut transportation times all around and how to provide the best service to the customer.

Although our data provided information only on the shipper of the unserviceable (retrograde pipeline) and not on the destinations for shipments of serviceables (order-and-ship pipeline), in our analysis of the G009 data we observed that at least 784 of the 862 recoverable assets received by contractors were shipped from Ogden to the contractor, having first been shipped from various bases to Ogden. Of the remaining assets, 33 were shipped directly from

bases and 45 had other identifiers.⁶ Assets that were shipped from Ogden spent an average of 36 days in transit to the contractor. Those shipped directly from bases took an average of 62 days to reach a contractor. (Those with other identifiers showed average shipment times of 68 days.) Again, because we did not have information on times for shipments of assets from bases to Ogden, we conclude that average retrograde transportation times for unserviceables shipped twice—from the base to Ogden and from Ogden to the contractor—are, on average, greater than 36 days, perhaps substantially greater. To enhance depot-level contractor component repair responsiveness, these transportation times need to be reduced substantially. One means of doing so is to ship unserviceables directly from bases to contractors, and to manage these pipelines carefully.

When we raised the idea of direct shipments of both serviceables and repairables between bases and the contractor, we received varying responses. Contractors argued that their costs could increase because of the extra handling in preparing shipping labels nearer to actual shipment dates. Furthermore, they pointed out that holding unserviceables at the contractor's facility rather than the ALC would add to the government's costs because of storage requirements. Ogden also noted that bases would need to know where to send unserviceables, something they can avoid if all shipments go to the managing ALC first. Distribution from the contractor to a base would require information beyond what is currently made available to the contractor, such as the intended destination's address for a particular NSN.

Considering all of the above, it strikes us that the advantages of a more responsive system that relies on more direct shipments between bases and the contractor clearly outweigh the startup complications implied by such a policy change, especially for scarce items. Since the government already pays for some amount of warehouse space at the contractor's facility, the ALC has some flexibility in determining what is stored on site. Projected repair requirements could guide the ALC in deciding whether unserviceables should go directly to a contractor or to the ALC, since the government may choose not to repair everything. Extra handling as a result of having to create shipping labels near the end of the repair process rather than the beginning must be balanced against the rather significant advantage of saving both time and spares investment requirements.

⁶For example, one unknown shipper was coded as "XXXXXX."

The current system does provide for the ALC to issue amended shipping instructions that authorize a change in the shipping destination from the one named in the contract. However, such instructions are used by exception, usually in the event of a MICAP. To ship items routinely from the contractor's door to a particular base would require instructions from the ALC on where to ship the asset. The decision should be predicated on which base would benefit the most from the receipt of the next available asset.

One possibility for making the decision on where to send the next available serviceable is to use DRIVE. According to a study conducted by the Logistics Management Institute, distribution decisions made by DRIVE outperformed both the item manager and the priority transportation set of procedures used by DoD, the Uniform Materiel Movement and Issue Priority System (UMMIPS) (Culosi and Eichorn, 1993). Ogden is using a desktop version of DRIVE for the distribution of some of the serviceables it ships to bases, namely, two-level maintenance F-16 avionics and lean logistics assets. This approach alleviates much of the overhead implied by real-time decisions and could account explicitly for worldwide aircraft availability goals. Changes to information systems and policies affecting direct retrograde and distribution of assets would be needed. At a minimum, contracts would have to change to accommodate these new requirements. The net benefit of improvements to responsiveness implicit in these changes is likely to be well worth the trouble.

Add Responsiveness Incentives to Repair Contracts

Much more than organic repair facilities, contractual sources of repair represent a so-called black box to the Air Force: Unserviceables go in and serviceables come out, but the processes that go on are poorly understood and hard to affect directly. To make the contractor's processes more responsive, one means open to the government is to motivate him to move in this direction through the use of explicit contractual incentives. Examples of such incentives might include:

- Specification of dramatically reduced repair flow times patterned after the Coronet Deuce experience
- Specification of repair flow times that include AWP and other delay times not due to any failure on the government's part
- Graduated monetary awards for achieving flow times shorter than specified, i.e., larger awards for greater differences
- Graduated penalties for failing to meet specified repair flow times, i.e., larger penalties for greater tardiness

- Monetary bonuses for meeting the specified flow times for sustained periods (a year, say)
- Sharing with the contractor the savings resulting from any cost-reduction initiatives the contractor may implement.

This set of examples includes emphasis on both cost and repair flow time. The second of these initiatives, including AWP and other delay times in the measured repair flow time, may appear troublesome at first glance, but we think it captures the spirit of getting the contractor involved, motivating him to accept the responsibility for achieving better performance by measuring his performance on all fronts, so to speak. By including delay times in the measured flow time, the Air Force may induce much more aggressive participation by the contractor in determining his own rewards. It seems an idea worth trying.

Most component repair contracts offer few incentives for the contractor to reduce repair flow times. During our interviews when we posed the question of including explicit responsiveness incentives in a contract and engaging the contractor to solve this problem, concern was raised about the potential cost effect. Hence, our brief list of examples includes cost as well as performance incentives. Again, it is important to keep in mind that shorter repair flow times will, in general, yield lower spares investment costs and may be attractive even when their achievement increases repair costs somewhat. Improved support to the combat force is an added incentive.

Use Responsiveness as a Contractor Selection Criterion

Although OSD has put formal public/private competition on hold for the time being, the Air Force could use repair flow times or broader measures of responsiveness as criteria for repair source selection, not only between organic and contractual sources, but also among alternative contractual sources. The mere statement of such a policy seems likely to have important motivational effects. The more emphasis the Air Force puts on repair responsiveness, the more pervasive it will become in repair management, and contractual repair is no exception.

Extend the Use of DRIVE to Contractual Repair

Assuming some changes to certain data systems and policies, DRIVE could support three types of contractual repair-related decisions: (1) where to ship the next serviceable asset that emerges from the repair process, discussed earlier; (2)

how much and what mix of workload to allocate to contractual sources of repair; and (3) given choices of what to induct into repair next, which asset the contractor should choose. The use of DRIVE for these purposes could enhance the relevance of contractual component repair.

DRIVE is a repair management decision support system that has been shown to improve the responsiveness of organic component repair.⁷ It works on the principle that if the organic shop has a choice of what asset to repair next, it should choose the one that provides the best benefit to its customer base, taking explicit account of the repair cost; it also helps in deciding which customer should receive the next serviceable. A successful demonstration of DRIVE at the Ogden Air Logistics Center on F-16 avionics organic workload occurred in 1987. Since then the Air Force has developed two versions of the system: a mainframe-based system located at HQ AFMC, called D087, and a PC-based system located at HQ AFMC and the ALCs, called *Desktop DRIVE*. Several ALCs reportedly run DRIVE to estimate repair workload requirements over short planning horizons, typically three to four weeks; Ogden uses it to determine quarterly workloads as well.

The suggestions made about DRIVE derive from observing its use in organic depot-level operations and considering its application to contractual repair responsiveness. It makes sense to use it to prioritize repairs whenever (a) there are choices that are not heavily constrained by setup costs or unserviceable availability, and (b) the shop has scope of repair. It always applies to the decision of where to ship the serviceable asset that emerges from repair. So long as the contractor has no compelling reason to repair components in a particular sequence, he should be willing to repair next the asset that the Air Force values the most in terms of its aircraft availability goals, and DRIVE provides that information.⁸

It should be noted that setting priorities is no substitute for better material support policies. If the contractor's flow times are dictated by the repair process itself, either because of job-routed SRUs or unavailable GFM, then setting repair priorities will help, but not as much as otherwise. In this role, DRIVE could help the Air Force implement proactive SRU repair and improve material support of the contract repair process, barring any prohibition against stockage of repair parts in anticipation of demand.

⁷DRIVE is well documented (Abell et al., 1992; Miller and Abell, 1992).

⁸Repair process data, such as standard hours, could be based on the Air Force's own estimates. Using this source of information would avoid having to gain access to potentially sensitive information. These estimates would apply to processes that are similar to those the Air Force uses and could be corroborated through normal inspections conducted before the contract award.

As an aside, one could imagine an incentive contract that rewarded contractors for following Air Force priorities. Such an approach has two obvious benefits: It motivates the contractor to get involved with solving material support problems, and it enables the Air Force to affect the contractor's operations from the outside. The rewards need not be financial. In today's world, just knowing the customer's preference is often enough to influence a supplier.

To make the above suggestion work, the ALC and the contractor would need to be able to pass lists of items back and forth daily. The ALC could assign addresses to the next several assets of each type the contractor repairs and update the list daily, and could also communicate repair priorities to the contractor for each production period.

Improve Requirements Determination and Workload Allocation to Reduce Uncertainty

For all of the reasons discussed in Section 3, uncertainty in requirements costs the Air Force money. The more stability contractual repair management can bring to the process underlying the issuance of funded delivery orders to the contractor, the less repairs are likely to cost.

Unless the contractor has the same equipment available to him as the government does, such as the automated test stands used for the repair of several types of avionics components, workloads that fluctuate create churn in the contractor's work force, particularly with dedicated repair facilities. The costs of these inefficiencies show up when the contractor must apply fixed costs over less workload, lay off his workers, then bring them back or replace them and train new ones, or pay overtime to meet surges and shift to less-profitable tasks such as cleaning up the work station area until work materializes again.

For the contractor to avoid losses, he must pass all of his costs on to the Air Force. In the process of negotiating a component repair contract, uncertainty in the repair requirement over the agreement's lifetime affects the unit repair cost. If the requirement is uncertain, the contractor may commit more resources than will ultimately be required or have to acquire additional resources to meet unanticipated demands. Either situation is more costly to him than having a stable workload and allocating resources to it accordingly. Thus, he must negotiate a higher unit price for uncertain requirements than he would need to for known, stable requirements, passing on the higher costs to the Air Force. In addition, unstable workloads compound the difficulty of the material support problem.

Many events may intervene between the requirements determination process itself and the issuance of the delivery order, most notably changes in funding levels and funding allocations. As we discussed earlier, contractors are vulnerable to the vagaries of funding as well as unserviceable generations and repair parts demands.

We see two approaches to reducing the uncertainty with which the contractor is forced to cope. The first is to stabilize repair requirements and funding; the second, for split items, is to allocate a fixed, stable portion of the workload to contractors and the remaining, more uncertain part of the workload to organic repair facilities.

Stabilize Repair Requirements and Funding

The primary way to address this problem is to attack the drivers leading to quantity and mix uncertainties to the extent possible. We described in the last section how repair requirements are estimated across the contract's lifetime and on a quarterly basis. At least two types of uncertainty exist that make these estimations difficult: uncertain scenarios and uncertain demands. Traditional methods rely on trends and the judgment of the IM and PMS to cope with both. A different approach would be to account for these two sources of uncertainty separately, by estimating expected demands in a probabilistic sense against specified aircraft availability goals. To the extent that estimation methods narrow uncertainties from both sources, the ALC strengthens its negotiating position. If these uncertainties were diminished, it seems likely that the Air Force would not be forced to use such small minimum quantities in its contract negotiations; thus, it may be possible to achieve lower unit costs.

Allocate the Uncertain Portion of the Workload for Split Items to Organic Sources of Repair

The basis for this idea is that organic repair shops typically have greater scope of repair than contractors. Scope of repair applies in two ways. Some pieces of test equipment are designed to test a wide variety of components. This is the type of test equipment often purchased by the Air Force, especially for avionics repair. The greater the scope of repair of the repair equipment, the more flexible the shop is in dealing with uncertain unserviceable generations. Thus, if less than the expected number of unserviceables of one item are generated, chances are that there will be more of another type.

The second source of scope of repair is simply in the range of items repaired. The greater the range of items, the greater the scope of repair. Large scope of repair implies the ability to choose from among a larger number of items.

Both kinds of scope of repair tend to be greater in organic repair shops than at contractors' facilities. This observation suggests the hypothesis, not examined in this research, that the best overall cost-effectiveness for split items might be achieved through the strategy of allocating a fixed portion of the workload to contractual sources and the remainder to organic sources, the opposite of the traditional approach. The advantage of doing so would be to reduce the costs of what cannot be managed closely (the contractor's operation) and to buy flexibility from the more robust source of repair that can be managed closely (the organic operation). This hypothesis could easily be tested, and could lead to substantial reductions in unit repair costs (organic costs would also need to be monitored). DRIVE could be helpful in this repair allocation decisionmaking.

We estimate that about two-thirds of the value of the split-item workload constitutes the underlying stable repair demand (see Figure 4.2).⁹ These estimates come from data during the 1980–1988 period; we counted as "split" anything repaired by both the private and public sectors within a given year.




Figure 4.2—Expenditures by Source Type: 1980–1988

⁹Data source: Weapon System Cost Retrieval System.

On the face of it, the above approach may seem counterintuitive. After all, if the organic shop has a lower unit repair cost than a contractor and the Air Force made a large investment in capability and capacity, then why shouldn't the government favor the less expensive source of repair? The answer lies in the joint consideration of these two sources of repair. If, overall, the contractor's unit repair costs and responsiveness improve more than the Air Force's cost and responsiveness degrade, then for good business reasons the Air Force ought to allocate the stable portion of the workload to the contractor and the uncertain portion to the organic shop.

5. Summary of Recommendations

In this section we summarize the recommendations described earlier. To begin, some office within AFMC and each ALC should have sufficient oversight authority to initiate policy changes that will lead to greater contractor responsiveness. That office represents the buyer of contractual services. It needs to have the authority and incentives to initiate improvements to contractual repair policy that will enhance repair responsiveness and increase its cost effectiveness.

Emphasizing the value of contractor responsiveness throughout all aspects of contractual component repair management is an important step toward achieving shorter flow times. Developing an objective, accurate means of measuring contractor performance provides many benefits, not just in awarding contracts, but also in their day-to-day management. Incorporating repair flow times, both with and without AWP times, as one dimension of depot-level costs gives decisionmakers a tool they need. They cannot just cut costs without taking responsiveness into consideration and the only way to do that is to measure it. We suggest that modest changes to G009 reporting requirements would provide such information on an important class of contracts. The emphasis on responsiveness must also extend to contractual incentives. Requirements contracts fail to provide the right ones and result in higher total costs when tradeoffs between repair cycle times and spares investments are explicitly considered.

Rethinking material support policies addresses a critical obstacle to enhanced responsiveness. Since the contractor relies on the MILSTRIP system for requisitioning repair parts and generally does not lay in adequate stocks of them, the Air Force does not count the time that elapses while the contractor waits for the delivery of government-furnished material on order. Thus, a contractor's performance is not affected by AWP time and he has little motivation to raise material support issues with the Air Force. Given the fact that the government currently manages many of these parts, the approach is fair. But the Air Force has opportunities to reduce total repair flow times through material support policies that lead to fewer AWP incidents. A policy that calls for more rotatable spares or on-hand serviceable SRUs could reduce the time an LRU must wait during the repair process, thus improving responsiveness.

Improving the process by which the ALCs generate contract repair requirements used in negotiating repair contracts would help reduce the amount of uncertainty involved with anticipated repair quantities, so necessary to reduce unit repair costs. Once the ALCs begin measuring repair flow times, they can take the next step in determining the value of better responsiveness in terms of total costs to the depot. Many people at the ALCs said that great uncertainty over projected repair requirements undercuts the government's ability to negotiate more favorable terms. Linking the repair estimation process to aircraft availability goals could help in projecting better repair mixes and quantities.

Eliminating secondary shipments to the fullest extent possible would vastly reduce transportation times and increase depot-level responsiveness to the customer. Retrograde and resupply delays increase with every unnecessary shipment of an item from a base or contractor to an ALC. Each time an item passes through an ALC, pipeline times increase. Moreover, the ALC can distribute items more effectively if it updates an item's destination as close to the end of the repair process as possible. AFMC could require that all contracts conform to this policy. Such an approach would lead to immediate benefits in getting assets to the customer faster. Air Force bases need to know where to send their unserviceables and the contractor needs to know where to ship the serviceables. Both decisions could be made in near-real-time, provided some changes in key information systems and transportation policies were made.

Evaluating the effectiveness and total costs of moving to a policy that allocates the uncertainty in predicting repairable generations to the organic source of repair for split items could offer new opportunities for reducing contractual repair uncertainties and repair costs.

Extending the use of Desktop DRIVE to contractual component repair for repair requirements determination, recoverable asset management, and serviceable asset allocation may have significant payoff. This methodology could support all of these processes. In situations where there is sufficient scope of repair, DRIVE's application can be extended to repair prioritization. In situations where the contractor repairs both LRUs and their indentured recoverable SRUs, DRIVE's application can be extended to proactive SRU repair.

Finally, we note that these policy recommendations imply various changes to current practices or different ways of considering processes. Importantly, all share a common overarching goal, that is an Air Force logistics support system that is more responsive and robust to a combat force in a continually changing threat environment. The significant contribution that contractual repair makes to

depot-level logistics support means that as the Air Force improves its organic repair operations, it has an important counterpart in its private sector source of repair to bring along as well.

Appendix

A. Contractor Repair Performance Information Systems

In 1992, the Office of Federal Procurement Policy (OFPP), Office of Management and Budget, directed all executive departments and establishments, including DoD and the services, to establish "requirements for evaluating contractor performance and for using past performance information in the contractor selection process" (OFPP, 1992). It directed all services to develop qualitative and quantitative measures that link a contractor's past performance to new contracts valued at \$100,000 or more.

AFMC is now in the process of interpreting this directive within the context of repair contracts. AFMC has three past performance data systems in varying stages of development that apply to replenishment spares acquisition contracts or weapon system acquisition programs. They are:

- The AFMC Blue Ribbon Contractor Program, used for contracts with three or more contract line items with a combined value of \$50,000 or more within 36 months. It contains information about on-time deliveries, quality, and overall performance. Contractors that qualify as Blue Ribbon Contractors may negotiate prices up to 10 or 20 percent higher than the lowest priced responsible offerer, not to exceed a given dollar threshold.
- The AFMC Vendor Rating System (VRS), used for contracts worth \$10,000 or more. The system contains information on contractor delivery and quality performance. It measures tardy deliveries, as well as the proportion of defective units compared to total units delivered. The planned date of availability for VRS was July 1994.
- The AFMC Contractor Performance Assessment Reporting System (CPARS), used for contracts valued at more than \$5 million. AFMC reportedly would like to modify this system for repair contracts. The government currently uses this information to make source-selection decisions for acquisition programs. It monitors overall performance, schedule performance, cost variance, product assurance, independent test and evaluation programs, logistics support, responsiveness, and subcontract management. These reports are tightly controlled and exist only in hard-copy form and in only a few locations.

As AFMC considers how it will meet these guidelines to extend past-performance considerations to repair contracts, we suggest that the command consider building a crosswalk between the G009, which potentially provides a good source of contractor responsiveness data (Air Force Materiel Command, 1993a) along with a data source for AWP information, and the eventual past-performance data system. Developing such a capability would implicitly enable the ALC to monitor responsiveness day-to-day as well.

B. Statistical Limitations of Responsiveness Data

In this appendix we discuss the statistical implications of responsiveness data that are reported by national stock number, but not serial number. We take the simplest example of one NSN, two assets, and two events in the repair process: arrival at the contractor's facility or an asset's receipt and its induction into repair.

Suppose the following information were reported in the G009: (1) A contractor records the receipt of two items of the same NSN, arriving one week apart from one another; (2) one week later, after this second arrival, he inducts one of the assets into repair and one week later, inducts another asset. Without serial number reporting, we cannot say whether the first asset that was inducted was the first or second one that arrived. If the first asset inducted was the first one that arrived, then the flow times for induction for these assets would be two weeks each. If, instead, the first asset the contractor inducted was actually the second one he received, then the respective times would be one and three weeks. Taken as a pair, their average flow times would be the same—two weeks. However, the variance in times in the first case is smaller than in the second case. This example illustrates the important principle that in the case of a balanced set of events, i.e., when the same number of receipts and inductions occur, the mean flow time is the same regardless of the true sequence of events, and a minimal variance will generate when a "first in, first out" assumption is applied.

In the case in which the contractor has not inducted all the assets he received, we have some uncertainty as to what portion of the received items are in the inducted pool. Therefore, the mean flow times computed for these assets are estimates of the true population mean (which we can compute exactly once all the assets are inducted).

The statistical implications of working with data without serial number identification require assumptions about the actual sequence of events. Consequently, our findings should be viewed only as indicative of responsiveness. We hope they also communicate the potential benefit of having visibility over actual repair flow times.

C. Contractor Responsiveness Data Analysis

The ALCs capture contractor responsiveness information in two ways: (1) through information systems that indicate asset balances in general terms, i.e., the number of assets sent to a contractor, received by a contractor, and delivered from a contractor; and (2) repair flow times as reported by the contractor himself. Independent verification of contractual repair flow times, i.e., those segments over which the contractor has control, is more difficult to obtain.

Ogden recommended two data systems as potential sources of detailed repair flow times: the Contract Depot Maintenance Production and Costs System (G072D) and the Government-Furnished Material Transaction Reporting System (G009). The G072D records repair delivery order quantities by quarter for a wide array of contracts. The G009 data system supplies information on contracts that use government-furnished material (repair parts that are purchased and managed by the government) of sufficient dollar value to merit the cost of requiring a contractor to provide information on a timely basis. Financial Management uses G009 to track government-furnished material inventories as required by the FARs, and G072D to account for work completed or still outstanding.

We were interested in high-resolution information that would allow us to compute repair flow times in units of days. Of the two data systems, G072D and G009, the latter provided more detailed information. G072D does not contain data on item inductions into the repair process at the contractor's facility or serviceable shipments by the contractor, but it does provide information on negotiated repair flow times and repair costs. G009 records repair transactions by calendar date on asset receipt, induction, production, and shipment, along with condemnations and GFM-specific occurrences. Neither system supplied adequate responsiveness data per se, but in combination they allowed us to infer something about timeliness.

Because G009 does not track assets by serial number, one cannot accurately infer many important statistics on timeliness. Without asset-specific identification, one can only know average repair flow times accurately and this only when all repair actions that were initiated have been completed. Measuring repair flow time, including the time when the asset being repaired is awaiting repair parts

requires the ability to follow a particular asset from the beginning of the repair process to the end. Lacking the information needed to distinguish one asset from another, we assumed that items that first entered the process were the first to emerge. This approach produces a conservative estimate of the standard deviation, and the same mean as the actual population of events in the sample. Without serial-number identification, we cannot know what actually happened asset by asset. We discussed the reasonableness and implications of this assumption in Appendix B.

Data quality problems with G009 restricted us to only those contracts that began reporting in January 1993 or later. The system had undergone a significant upgrade and over a period of several months allowed corrupted information to enter. On examining data from several contracts posted earlier than 1993, we found evidence indicative of those problems (e.g., the number of assets repaired exceeded the number inducted).

Finally, our original sample contained information on both F-16 and non-F-16 items. The F-16 items, making up over 80 percent of our sample, fell into two groups, those with 30-day negotiated flow times and those with 45-day NFTs. The flow times implied by G009 represent actual calendar time. Although we note the fact that AWP events may partially account for why repair flow times exceed their contractual timeliness requirements, we were not able to verify the extent to which these events affect overall responsiveness.

Bibliography

- Abell, John B., Grace M. Carter, Karen E. Isaacson, and Thomas F. Lippiatt, *Estimating Requirements for Aircraft Recoverable Spares and Depot Repair*, Santa Monica, California: RAND, R-4210-AF, 1993.
- Abell, John B., L. W. Miller, C. E. Neumann, and J. Payne, *DRIVE (Distribution and Repair in Variable Environments): Enhancing the Responsiveness of Depot Repair*, Santa Monica, California: RAND, R-3888-AF, 1992.
- Air Force Materiel Command, *Air Force Materiel Command Guide on Integrated Product Development: A Guide for Understanding and Implementing IPD Throughout AFMC*, Wright-Patterson AFB, Ohio: Headquarters, AFMC, 25 May 1993b.
- Air Force Materiel Command, *General Characteristics of Organic and Contract Depot Maintenance Workloads*, Wright-Patterson AFB, Ohio: Headquarters, AFMC, AFMC Regulation 500-26, 1 January 1994.
- Air Force Materiel Command, *Government-Furnished Material and End Item Transaction Reporting System (G009)*, Wright-Patterson AFB, Ohio: Headquarters, AFMC, AFMC Manual 66-266, Vol. I, 15 January 1993a.
- Comptroller of the Department of Defense, *Defense Business Operations Fund: Improvement Plan*, Washington, D.C.: Department of Defense, September 1993.
- Camm, Frank, David Dreyfuss, and Karl Hoffmayer, "The Role of the DoD Cost Comparability Handbook in Air Force Public/Private Competitions," Santa Monica, California: RAND, unpublished research, 1994.
- Cohen, Irving K., John B. Abell, and T. F. Lippiatt, *Coupling Logistics to Operations to Meet Uncertainty and the Threat (CLOUT): An Overview*, Santa Monica, California: RAND, R-3979-AF, 1991.
- Culosi, Salvatore J., and Frank L. Eichorn, *A Comparison of Two Systems for Distributing Spare Parts*, Bethesda, Maryland: LMI, AF201R1, March 1993.
- Donohue, George L., and Marygail Brauner, *Revamping the Infrastructure That Supports Military Systems*, Santa Monica, California: RAND, IP-140-AF, 1993.
- Grier, Peter, "The Other Industrial Base," *AIR FORCE Magazine*, July 1992, pp. 50-52.
- Jane's All the World Aircraft*, "Lockheed (General Dynamics) F-16 Fighting Falcon," Surrey, United Kingdom: Jane's Information Group, 1993-1994.

Kingsbury, Nancy R., *Depot Maintenance: Issues in Management and Restructuring to Support a Downsized Military*, Washington, D.C.: General Accounting Office, GAO/T-NSIAD-92-23, March 26, 1992.

Miller, Louis W., and John B. Abell, *DRIVE (Distribution and Repair in Variable Environments): Design and Operation of the Ogden Prototype*, Santa Monica, California: RAND, R-4158-AF, 1992.

Office of the Assistant Secretary of Defense, internal memorandum, "Policy Change Request (Buyer-Seller Relationships)," Washington, D.C., August 7, 1990.

Office of Federal Acquisition Policy, *Title 48, Code of Federal Regulations*, Part 45.505.1, Washington, D.C.: Government Printing Office, 1993.

Office of Federal Procurement Policy, Office of Management and Budget, *Past Performance Information*, Policy Letter No. 92-5, received December 30, 1992, Washington, D.C.: no date.

Office of Federal Procurement Policy, Office of Management and Budget, *Past Performance Information Systems*, Washington, D.C.: October 1993.